

1

Introduction

1.1 Motivation

Since the 1970s, a series of significant changes has taken place in the world of finance. One key year was 1973, when currencies began to be traded in financial markets and their values determined by the foreign exchange market, a financial market active 24 hours a day all over the world. During that same year, Black and Scholes [18] published the first paper that presented a rational option-pricing formula.

Since that time, the volume of foreign exchange trading has been growing at an impressive rate. The transaction volume in 1995 was 80 times what it was in 1973. An even more impressive growth has taken place in the field of derivative products. The total value of financial derivative market contracts issued in 1996 was 35 trillion US dollars. Contracts totaling approximately 25 trillion USD were negotiated in the over-the-counter market (i.e., directly between firms or financial institutions), and the rest (approximately 10 trillion USD) in specialized exchanges that deal only in derivative contracts. Today, financial markets facilitate the trading of huge amounts of money, assets, and goods in a competitive global environment.

A second revolution began in the 1980s when electronic trading, already a part of the environment of the major stock exchanges, was adapted to the foreign exchange market. The electronic storing of data relating to financial contracts – or to prices at which traders are willing to buy (bid quotes) or sell (ask quotes) a financial asset – was put in place at about the same time that electronic trading became widespread. One result is that today a huge amount of electronically stored financial data is readily available. These data are characterized by the property of being high-frequency data – the average time delay between two records can be as short as a few seconds. The enormous expansion of financial markets requires strong investments in money and

human resources to achieve reliable quantification and minimization of risk for the financial institutions involved.

1.2 Pioneering approaches

In this book we discuss the application to financial markets of such concepts as power-law distributions, correlations, scaling, unpredictable time series, and random processes. During the past 30 years, physicists have achieved important results in the field of phase transitions, statistical mechanics, nonlinear dynamics, and disordered systems. In these fields, power laws, scaling, and unpredictable (stochastic or deterministic) time series are present and the current interpretation of the underlying physics is often obtained using these concepts.

With this background in mind, it may surprise scholars trained in the natural sciences to learn that the first use of a power-law distribution – and the first mathematical formalization of a random walk – took place in the social sciences. Almost exactly 100 years ago, the Italian social economist Pareto investigated the statistical character of the wealth of individuals in a stable economy by modeling them using the distribution

$$y \sim x^{-\nu}, \tag{1.1}$$

where y is the number of people having income x or greater than x and ν is an exponent that Pareto estimated to be 1.5 [132]. Pareto noticed that his result was quite general and applicable to nations ‘as different as those of England, of Ireland, of Germany, of the Italian cities, and even of Peru’.

It should be fully appreciated that the concept of a power-law distribution is counterintuitive, because it may lack any characteristic scale. This property prevented the use of power-law distributions in the natural sciences until the recent emergence of new paradigms (i) in probability theory, thanks to the work of Lévy [92] and thanks to the application of power-law distributions to several problems pursued by Mandelbrot [103]; and (ii) in the study of phase transitions, which introduced the concepts of scaling for thermodynamic functions and correlation functions [147].

Another concept ubiquitous in the natural sciences is the random walk. The first theoretical description of a random walk in the natural sciences was performed in 1905 by Einstein [48] in his famous paper dealing with the determination of the Avogadro number. In subsequent years, the mathematics of the random walk was made more rigorous by Wiener [158], and

now the random walk concept has spread across almost all research areas in the natural sciences.

The first formalization of a random walk was not in a publication by Einstein, but in a doctoral thesis by Bachelier [8]. Bachelier, a French mathematician, presented his thesis to the faculty of sciences at the Academy of Paris on 29 March 1900, for the degree of *Docteur en Sciences Mathématiques*. His advisor was Poincaré, one of the greatest mathematicians of his time. The thesis, entitled *Théorie de la spéculation*, is surprising in several respects. It deals with the pricing of options in speculative markets, an activity that today is extremely important in financial markets where derivative securities – those whose value depends on the values of other more basic underlying variables – are regularly traded on many different exchanges. To complete this task, Bachelier determined the probability of price changes by writing down what is now called the Chapman–Kolmogorov equation and recognizing that what is now called a Wiener process satisfies the diffusion equation (this point was rediscovered by Einstein in his 1905 paper on Brownian motion). Retrospectively analyzed, Bachelier’s thesis lacks rigor in some of its mathematical and economic points. Specifically, the determination of a Gaussian distribution for the price changes was – mathematically speaking – not sufficiently motivated. On the economic side, Bachelier investigated price changes, whereas economists are mainly dealing with changes in the logarithm of price. However, these limitations do not diminish the value of Bachelier’s pioneering work.

To put Bachelier’s work into perspective, the Black & Scholes option-pricing model – considered the milestone in option-pricing theory – was published in 1973, almost three-quarters of a century after the publication of his thesis. Moreover, theorists and practitioners are aware that the Black & Scholes model needs correction in its application, meaning that the problem of which stochastic process describes the changes in the logarithm of prices in a financial market is still an open one.

The problem of the distribution of price changes has been considered by several authors since the 1950s, which was the period when mathematicians began to show interest in the modeling of stock market prices. Bachelier’s original proposal of Gaussian distributed price changes was soon replaced by a model in which stock prices are log-normal distributed, i.e., stock prices are performing a geometric Brownian motion. In a geometric Brownian motion, the differences of the logarithms of prices are Gaussian distributed. This model is known to provide only a first approximation of what is observed in real data. For this reason, a number of alternative models have been proposed with the aim of explaining

- (i) the empirical evidence that the tails of measured distributions are fatter than expected for a geometric Brownian motion; and
- (ii) the time fluctuations of the second moment of price changes.

Among the alternative models proposed, ‘the most revolutionary development in the theory of speculative prices since Bachelier’s initial work’ [38], is Mandelbrot’s hypothesis that price changes follow a Lévy stable distribution [102]. Lévy stable processes are stochastic processes obeying a generalized central limit theorem. By obeying a generalized form of the central limit theorem, they have a number of interesting properties. They are stable (as are the more common Gaussian processes) – i.e., the sum of two independent stochastic processes x_1 and x_2 characterized by the same Lévy distribution of index α is itself a stochastic process characterized by a Lévy distribution of the same index. The shape of the distribution is maintained (is stable) by summing up independent identically distributed Lévy stable random variables.

As we shall see, Lévy stable processes define a basin of attraction in the functional space of probability density functions. The sum of independent identically distributed stochastic processes $S_n \equiv \sum_{i=1}^n x_i$ characterized by a probability density function with power-law tails,

$$P(x) \sim x^{-(1+\alpha)}, \tag{1.2}$$

will converge, in probability, to a Lévy stable stochastic process of index α when n tends to infinity [66].

This property tells us that the distribution of a Lévy stable process is a power-law distribution for large values of the stochastic variable x . The fact that power-law distributions may lack a typical scale is reflected in Lévy stable processes by the property that the variance of Lévy stable processes is infinite for $\alpha < 2$. Stochastic processes with infinite variance, although well defined mathematically, are extremely difficult to use and, moreover, raise fundamental questions when applied to real systems. For example, in physical systems the second moment is often related to the system temperature, so infinite variances imply an infinite (or undefined) temperature. In financial systems, an infinite variance would complicate the important task of risk estimation.

1.3 The chaos approach

A widely accepted belief in financial theory is that time series of asset prices are unpredictable. This belief is the cornerstone of the description of price

dynamics as stochastic processes. Since the 1980s it has been recognized in the physical sciences that unpredictable time series and stochastic processes are not synonymous. Specifically, chaos theory has shown that unpredictable time series can arise from deterministic nonlinear systems. The results obtained in the study of physical and biological systems triggered an interest in economic systems, and theoretical and empirical studies have investigated whether the time evolution of asset prices in financial markets might indeed be due to underlying nonlinear deterministic dynamics of a (limited) number of variables.

One of the goals of researchers studying financial markets with the tools of nonlinear dynamics has been to reconstruct the (hypothetical) strange attractor present in the chaotic time evolution and to measure its dimension d . The reconstruction of the underlying attractor and its dimension d is not an easy task. The more reliable estimation of d is the inequality $d > 6$. For chaotic systems with $d > 3$, it is rather difficult to distinguish between a chaotic time evolution and a random process, especially if the underlying deterministic dynamics are unknown. Hence, from an empirical point of view, it is quite unlikely that it will be possible to discriminate between the random and the chaotic hypotheses.

Although it cannot be ruled out that financial markets follow chaotic dynamics, we choose to work within a paradigm that asserts price dynamics are stochastic processes. Our choice is motivated by the observation that the time evolution of an asset price depends on all the information affecting (or believed to be affecting) the investigated asset and it seems unlikely to us that all this information can be essentially described by a small number of nonlinear deterministic equations.

1.4 The present focus

Financial markets exhibit several of the properties that characterize complex systems. They are open systems in which many subunits interact nonlinearly in the presence of feedback. In financial markets, the governing rules are rather stable and the time evolution of the system is continuously monitored. It is now possible to develop models and to test their accuracy and predictive power using available data, since large databases exist even for high-frequency data.

One of the more active areas in finance is the pricing of derivative instruments. In the simplest case, an asset is described by a stochastic process and a derivative security (or contingent claim) is evaluated on the basis of the type of security and the value and statistical properties of the underlying

asset. This problem presents at least two different aspects: (i) ‘fundamental’ aspects, which are related to the nature of the random process of the asset, and (ii) ‘applied’ or ‘technical’ aspects, which are related to the solution of the option-pricing problem under the assumption that the underlying asset performs the proposed random process.

Recently, a growing number of physicists have attempted to analyze and model financial markets and, more generally, economic systems. The interest of this community in financial and economic systems has roots that date back to 1936, when Majorana wrote a pioneering paper on the essential analogy between statistical laws in physics and in the social sciences [101]. This unorthodox point of view was considered of marginal interest until recently. Indeed, prior to the 1990s, very few professional physicists did any research associated with social or economic systems. The exceptions included Kadanoff [76], Montroll [125], and a group of physical scientists at the Santa Fe Institute [5].

Since 1990, the physics research activity in this field has become less episodic and a research community has begun to emerge. New interdisciplinary journals have been published, conferences have been organized, and a set of potentially tractable scientific problems has been provisionally identified. The research activity of this group of physicists is complementary to the most traditional approaches of finance and mathematical finance. One characteristic difference is the emphasis that physicists put on the empirical analysis of economic data. Another is the background of theory and method in the field of statistical physics developed over the past 30 years that physicists bring to the subject. The concepts of scaling, universality, disordered frustrated systems, and self-organized systems might be helpful in the analysis and modeling of financial and economic systems. One argument that is sometimes raised at this point is that an empirical analysis performed on financial or economic data is not equivalent to the usual experimental investigation that takes place in physical sciences. In other words, it is impossible to perform large-scale experiments in economics and finance that could falsify any given theory.

We note that this limitation is not specific to economic and financial systems, but also affects such well developed areas of physics as astrophysics, atmospheric physics, and geophysics. Hence, in analogy to activity in these more established areas, we find that we are able to test and falsify any theories associated with the currently available sets of financial and economic data provided in the form of recorded files of financial and economic activity.

Among the important areas of physics research dealing with financial and economic systems, one concerns the complete statistical characterization of

the stochastic process of price changes of a financial asset. Several studies have been performed that focus on different aspects of the analyzed stochastic process, e.g., the shape of the distribution of price changes [22, 64, 67, 105, 111, 135], the temporal memory [35, 93, 95, 112], and the higher-order statistical properties [6, 31, 126]. This is still an active area, and attempts are ongoing to develop the most satisfactory stochastic model describing all the features encountered in empirical analyses. One important accomplishment in this area is an almost complete consensus concerning the finiteness of the second moment of price changes. This has been a longstanding problem in finance, and its resolution has come about because of the renewed interest in the empirical study of financial systems.

A second area concerns the development of a theoretical model that is able to encompass all the essential features of real financial markets. Several models have been proposed [10, 11, 23, 25, 29, 90, 91, 104, 117, 142, 146, 149–152], and some of the main properties of the stochastic dynamics of stock price are reproduced by these models as, for example, the leptokurtic ‘fat-tailed’ non-Gaussian shape of the distribution of price differences. Parallel attempts in the modeling of financial markets have been developed by economists [98–100].

Other areas that are undergoing intense investigations deal with the rational pricing of a derivative product when some of the canonical assumptions of the Black & Scholes model are relaxed [7, 21, 22] and with aspects of portfolio selection and its dynamical optimization [14, 62, 63, 116, 145]. A further area of research considers analogies and differences between price dynamics in a financial market and such physical processes as turbulence [64, 112, 113] and ecological systems [55, 135].

One common theme encountered in these research areas is the time correlation of a financial series. The detection of the presence of a higher-order correlation in price changes has motivated a reconsideration of some beliefs of what is termed ‘technical analysis’ [155].

In addition to the studies that analyze and model financial systems, there are studies of the income distribution of firms and studies of the statistical properties of their growth rates [2, 3, 148, 153]. The statistical properties of the economic performances of complex organizations such as universities or entire countries have also been investigated [89].

This brief presentation of some of the current efforts in this emerging discipline has only illustrative purposes and cannot be exhaustive. For a more complete overview, consider, for example, the proceedings of conferences dedicated to these topics [78, 88, 109].

2

Efficient market hypothesis

2.1 Concepts, paradigms, and variables

Financial markets are systems in which a large number of traders interact with one another and react to external information in order to determine the best price for a given item. The goods might be as different as animals, ore, equities, currencies, or bonds – or derivative products issued on those underlying financial goods. Some markets are localized in specific cities (e.g., New York, Tokyo, and London) while others (such as the foreign exchange market) are delocalized and accessible all over the world.

When one inspects a time series of the time evolution of the price, volume, and number of transactions of a financial product, one recognizes that the time evolution is unpredictable. At first sight, one might sense a curious paradox. An important time series, such as the price of a financial good, is essentially indistinguishable from a stochastic process. There are deep reasons for this kind of behavior, and in this chapter we will examine some of these.

2.2 Arbitrage

A key concept for the understanding of markets is the concept of arbitrage – the purchase and sale of the same or equivalent security in order to profit from price discrepancies. Two simple examples illustrate this concept. At a given time, 1 kg of oranges costs 0.60 euro in Naples and 0.50 USD in Miami. If the cost of transporting and storing 1 kg of oranges from Miami to Naples is 0.10 euro, by buying 100,000 kg of oranges in Miami and immediately selling them in Naples it is possible to realize a risk-free profit of

$$100,000[0.60 - (0.80 \times 0.50) - 0.10] = 10,000 \text{ euro.} \tag{2.1}$$

2.3 *Efficient market hypothesis* 9

Here it is assumed that the exchange rate between the US dollar and the euro is 0.80 at the time of the transaction.

This kind of arbitrage opportunity can also be observed in financial markets. Consider the following situation. A stock is traded in two different stock exchanges in two countries with different currencies, e.g., Milan and New York. The current price of a share of the stock is 9 USD in New York and 8 euro in Milan and the exchange rate between USD and euro is 0.80. By buying 1,000 shares of the stock in New York and selling them in Milan, the arbitrager makes a profit (apart from transaction costs) of

$$1,000 \left(\frac{8}{0.80} - 9 \right) = 1,000 \text{ USD.} \tag{2.2}$$

The presence of traders looking for arbitrage conditions contributes to a market’s ability to evolve the most rational price for a good. To see this, suppose that one has discovered an arbitrage opportunity. One will exploit it and, if one succeeds in making a profit, one will repeat the same action. In the above example, oranges are bought in Miami and sold in Naples. If this action is carried out repeatedly and systematically, the demand for oranges will increase in Miami and decrease in Naples. The net effect of this action will then be an increase in the price of oranges in Miami and a decrease in the price in Naples. After a period of time, the prices in both locations will become more ‘rational’, and thus will no longer provide arbitrage opportunities.

To summarize: (i) new arbitrage opportunities continually appear and are discovered in the markets but (ii) as soon as an arbitrage opportunity begins to be exploited, the system moves in a direction that gradually eliminates the arbitrage opportunity.

2.3 Efficient market hypothesis

Markets are complex systems that incorporate information about a given asset in the time series of its price. The most accepted paradigm among scholars in finance is that the market is highly efficient in the determination of the most rational price of the traded asset. The efficient market hypothesis was originally formulated in the 1960s [53]. A market is said to be efficient if all the available information is instantly processed when it reaches the market and it is immediately reflected in a new value of prices of the assets traded.

The theoretical motivation for the efficient market hypothesis has its roots in the pioneering work of Bachelier [8], who at the beginning of the twentieth

century proposed that the price of assets in a speculative market be described as a stochastic process. This work remained almost unknown until the 1950s, when empirical results [38] about the serial correlation of the rate of return showed that correlations on a short time scale are negligible and that the approximate behavior of return time series is indeed similar to uncorrelated random walks.

The efficient market hypothesis was formulated explicitly in 1965 by Samuelson [141], who showed mathematically that properly anticipated prices fluctuate randomly. Using the hypothesis of rational behavior and market efficiency, he was able to demonstrate how Y_{t+1} , the expected value of the price of a given asset at time $t + 1$, is related to the previous values of prices $Y_0, Y_1, \dots, Y_t$ through the relation

$$E\{Y_{t+1}|Y_0, Y_1, \dots, Y_t\} = Y_t. \tag{2.3}$$

Stochastic processes obeying the conditional probability given in Eq. (2.3) are called martingales (see Appendix B for a formal definition). The notion of a martingale is, intuitively, a probabilistic model of a ‘fair’ game. In gambler’s terms, the game is fair when gains and losses cancel, and the gambler’s expected future wealth coincides with the gambler’s present assets. The fair game conclusion about the price changes observed in a financial market is equivalent to the statement that there is no way of making a profit on an asset by simply using the recorded history of its price fluctuations. The conclusion of this ‘weak form’ of the efficient market hypothesis is then that price changes are unpredictable from the historical time series of those changes.

Since the 1960s, a great number of empirical investigations have been devoted to testing the efficient market hypothesis [54]. In the great majority of the empirical studies, the time correlation between price changes has been found to be negligibly small, supporting the efficient market hypothesis. However, it was shown in the 1980s that by using the information present in additional time series such as earnings/price ratios, dividend yields, and term-structure variables, it is possible to make predictions of the rate of return of a given asset on a long time scale, much longer than a month. Thus empirical observations have challenged the stricter form of the efficient market hypothesis.

Thus empirical observations and theoretical considerations show that price changes are difficult if not impossible to predict if one starts from the time series of price changes. In its strict form, an efficient market is an idealized system. In actual markets, residual inefficiencies are always present. Searching